arXiv:1606.01742v1 [cs.DB] 6 Jun 2016

Adaptive Distributed Top-k Query Processing

Claus Dabringer and Johann Eder *

Alpe Adria University Klagenfurt, Department of Informatics Systems
{Claus.Dabringer,Johann.Eder}@aau.at

Abstract. ADiT is an adaptive approach for processing distributed top-k queries over peer-to-peer networks optimizing both system load and query response time. This approach considers the size of the peer to peer network, the amount k of searched objects, the network capabilities of a connected peer, i.e. the transmission rate, the amount of objects stored on each peer, and the speed of a peer in processing a local top-k query. In extensive experiments with a variety of scenarios we could show that ADiT outperforms state of the art distributed query processing techniques.

Keywords: distributed query processing, top-K query, peer-to-peer databases, federated databases

1 Introduction

Top-k queries retrieve the k tuples of a query result which score best for a given objective function. Top-k queries help to overcome the problem of too large query results on one hand and too low recall, if the query is more constrained, and are therefore a promising technology for improving and accelerating search in for various data collections, e.g. for the search for suitable samples in biobanks [10], our main application area. Top-k queries are also popular for providing users ranked search results they are used from web search engines. Top-k queries, in particular, the optimization of top-k query processing for central databases received a lot of attention [3,5,8,9,18,19,20,21,23,26]. Optimizing top-k queries in distributed environment, in particular in highly distributed networks of federated or peer to peer databases still has significant research needs.

Current distributed top-k query processing approaches either focus on reducing the amount of transmitted queries [2,17] or on keeping the amount of transported objects low [4,6,12,24], or to reduce the communication costs[11]. However, both the transmitted objects and messages affect the *system effort* and *query response time* in a peer to peer system. Therefore, we introduce an adaptive distributed top-k (short ADiT) query processing approach considering both. To the best of our knowledge ADiT is the first approach using the amount

* The work reported here was supported by the Austrian Ministry of Science and Research within the program GENAU (project GATIB II) and within the project BBMRI.AT.

of messages *and* the amount of transmitted objects to measure *system effort* and *query response time*.

Processing a top-k query in a p2p network with horizontal partitioning involves sending a top-k query to each peer. The optimization problem now is to determine a proper k_p for each p of the peers, i.e. how many objects should be fetched from which peer. If this k_p is too large, it results in unnecessary computation at the peers' site and unnecessary traffic. If it is too low, it is necessary to send additional queries to the peers.

In our approach several parameters are used to calculate a proper k_p:

- size of the peer to peer network
- amount k of searched objects
- network capabilities of each peer, i.e. the transmission rate
- amount of objects stored on each connected peer
- speed of a peer, i.e. the searching performance of that peer

In the following we show the general architecture for processing top-k queries in a p2p environment and derive some heuristics based on the parameters outlined above. We describe the implementation of the ADiT approach and show in an extensive set of experiments the performance gains using this approach.

2 ADiT in General

Adaptive Distributed Top-K query processing (short ADiT) is able to process distributed top-k queries over horizontally partitioned data exactly. ADiT assumes a dynamic peer to peer network. Each peer has variable bandwidth capabilities and individual message costs. In contrast to other approaches [24,25] ADiT does not rely on caching techniques. Thus the performance is not dependent on stable data or on reoccurring queries.

The aim of ADiT is to achieve a low overall system effort as well as a fast query response time. The first parameter, the overall system effort is defined as sum all amounts of time of the peers needed for (1) sending requests to other peers in the network to obtain further objects, (2) searching objects and (3) transmitting objects. The second parameter is the query response time, the time elapsed between submitting a query and the return of the result. Formula 1 and formula 2 define the system effort, respectively the query response time where $MsgCount_i$ is the total amount of messages sent to peer P_i and n_i is the amount of objects retrieved from peer P_i. We use the following abbreviations throughout of this paper: N is the peer to peer network, Q is the top-k query, R is the queried relation, and P_i is a peer in the peer to peer system.

$$SE(N,Q,R) = \sum_{i=1}^{|P|} CCN.P_i, MsgC_i) + DBCosts(N.P_i, Q, R, n_i) + TransCosts(N.P_i, R, n_i) \quad (1)$$

$$QueryAnswerTime(N, Q, R) = max(CommCosts(N.P_i, MsgCount_i), \quad (2)$$
$$DBCosts(N.P_i, Q, R, n_i),$$
$$TransCosts(N.P_i, R, n_i))$$

The unit of system effort as well as of query response time is seconds. Thus it is needed to map the different costs to a time factor. Function 3 defines how sending $MsgCount$ requests to peer P is mapped to a time factor. The amount of incoming messages is multiplied with the constant costs that arise when establishing a connection to peer P. This gives the amount of time that is spent by sending $MsgCount$ messages to peer P.

$$CommCosts(P, MsgCount) = P_{MsgCosts} * MsgCount \quad (3)$$

Function 4 defines how retrieving n objects from relation R of peer P is mapped to a time factor. The transmission costs are influenced by the size of the object in relation R on peer P and by the transmission rate of peer P.

$$TransCosts(P, R, n) = \frac{(P_{R_{ObjectSize}} * n)}{P_{TransRate}} \quad (4)$$

The database costs ($DBCosts(N.P_i, Q, R, n)$) for searching the best n objects in relation R on peer P_i strongly depend on the top-k approach used on peer P_i, performance of the answering peer P_i, and the issued query Q, e.g. on the number of restrictions. ADiT assumes that each peer provides an estimate of the time needed to return the *top-k* objects for a query with m restrictions on a relation with size N. There is no assumption which procedure a peer uses to process top-k queries..

ADiT works iteratively and calculates a separate fetch size k'_p for each peer in each iteration. Then ADiT broadcasts the query Q *in parallel* and gathers the top-k'_p from each peer p. Then ADiT tries to publish objects and repeats if necessary.

There are two major possibilities for tuning: Choosing an appropriate fetch size k'_p for each peer in each iteration and avoiding to contact peers which cannot contribute to the result. For choosing the fetch size there are two extreme cases:

1. Setting $k'_p = 1$ for each peer leads to a minimal amount of *transmitted objects* but to a higher amount of *transmitted messages*.
2. Setting $k'_p = k$ for each peer leads to a minimal amount of *transmitted messages* but to a higher amount of *transmitted objects*.

In the rest of this paper we will focus on how to tune this basic distributed top-k query processing approach.

3 Heuristic Fetch Size Calculations

Analyzing a large number of queries varying the influencing factors [7] we developed two heuristics (basic and enhanced) for choosing a good fetch size k'_p for each individual peer p.

Basic Heuristics. The basic heuristics shown in equation 5 only uses the amount of relevant peers N_{Size} and the amount of searched objects k to derive a common fetch size f for all peers. The basic heuristics does not assume any particular data distribution. Thus it tries to retrieve an equal amount of objects from each peer. In case k is larger than N_{Size} the basic heuristics equally distributes k among the available peers. Otherwise the basic heuristics calculates the smallest multiple of k which is greater or equal than N_{Size} and equally distributes this amount among the available peers. The *consFactor* is used to increase the fetch size since it is unlikely that each peer will contribute the same number of objects. This increasing is used to fetch more objects and keep the number of iterations small. Our initial experiments showed that a *consFactor* of 2 leads to good results, e.g. few iterations and thus few messages exchanged in the p2p network. If the data is not distributed equally, *consFactor* should be chosen higher.

$$f = min(k, consFactor * \left\lceil \frac{N_{Size}}{k} \right\rceil * \frac{k}{N_{Size}})) \tag{5}$$

Enhanced Heuristics. The enhanced heuristics calculates the fetch size k'_p for each peer p *separately*. It uses additional parameters to adjust the fetch size for each peer properly:

- $ObjectsStored_p$: Amount of objects stored on peer p.
- $ObjectsStored_N$: Amount of objects stored in the peer to peer system N, i.e. $sum(ObjectsStored_p)$
- $Speed_p$: Query processing speed of peer p, e.g. a value between 1 and 10 where 1 is the slowest and 10 the fastest speed.
- $maxSpeed_N$: Maximum query processing speed of a peer in the peer to peer system N.
- $TransRate_p$: The transmission rate describing how fast the network connection of a certain peer is. This value is given in MBit per second.
- $maxTransRate_N$: Maximum transmission rate of a peer in the peer to peer system N.

The knowledge gathered during query processing iterations comprises the following parameters:

- $ObjectsRetrieved_p$: Amount of objects of peer p which have already been retrieved, initially 0.
- $ObjectsPublished_p$: Amount of objects of peer p which made it in the top-k answers, initially 0.
- $ObjPub_N$: Amount of objects returned to the user, initially 0.

All these parameters are used to calculate different weights which influence the enhanced heuristics. Applying the basic heuristics to the large test scenarios showed that the proposed fetch size should be treated as a lower limit. Therefore, the enhanced heuristics uses the different weights to *increase the fetch size* determined with the basic heuristics. To accomplish that the enhanced heuristics maps its weights to the interval of [1, 2]. This prevents from fetching fewer objects than the basic heuristics suggested. The enhanced heuristics assumes that all previous iterations can be used to reason about following iterations, e.g. it assumes that peers that contributed more objects in previous iterations will also contribute more objects in the following iterations. This assumption is reflected in weight w_{pF} which is defined in equation 6. The more objects a peer published compared to all other peers, the more objects are gathered from this peer *in the next iteration.*

$$w_{pF} = (1 + \frac{ObjectsPublished_p}{ObjPub_N}) \tag{6}$$

The enhanced heuristics tries to reduce the amount of fetched objects which are not needed. Thus it fetches more objects from peers where the ratio between fetched objects and published objects is high. Equation 7 shows the definition of weight w_{uF}.

$$w_{uF} = (1 + \frac{ObjectsPublished_p}{ObjectsRetrieved_p}) \tag{7}$$

The enhanced heuristics assumes that peers which store more objects will contribute more to the final answer. Thus it suggests to fetch more objects from larger peers. It uses equation 8 to incorporate that fact, namely weight w_{DBF}.

$$w_{DBF} = (1 + \frac{ObjectsStored_p}{ObjectsStored_N}) \tag{8}$$

Since it is cheap to ask a faster peer for more objects the enhanced heuristics defines w_{Speed} and $w_{TransRate}$. Equation 9 models the fact that more objects should be fetched from peers which are faster in searching their databases.

$$w_{Speed} = (1 + \frac{Speed_p}{maxSpeed_N}) \tag{9}$$

Equation 10 deals with the transmission of objects. It reflects that more objects should be fetched from peers which have a higher transmission rate.

$$w_{TransRate} = (1 + \frac{TransRate_p}{maxTransRate_N}) \tag{10}$$

The weights described in equations 6-10 are used by the enhanced heuristics to influence the basic heuristics. The weighted fetch size is determined with the heuristic function shown in equation 11.

$$k'_p = min(k - ObjPub_N, \lceil f * w_{pF} * w_{uF} * w_{DBF} * w_{Speed} * w_{TransRate} \rceil) \tag{11}$$

The upper bound for fetch size k'_p is obviously the amount of missing objects, namely $k - ObjPub_N$. The enhanced heuristics does not fetch more objects than the amount of missing objects from any of the peers in the peer to peer system.

3.1 ADiT Processing Iterations

ADiT processes a given distributed top-k query through a number of iterations. Each iteration is used to gather objects from the peers within the system to satisfy the distributed top-k query. In this section we focus on the relevant steps in each iteration. The pseudo-code in listing 1.1 shows how ADiT obtains the best k objects for a list of restrictions.

The variables used for storing the maximum remaining score ($maxRemScore$) and all fetched objects ($fetchedObjs$) are assumed to be globally visible to all threads during execution. They are depicted as in-out parameters in all pseudo-codes where they are used. The output produced by the $ADiT$-method is a sorted list of the k objects which score best among all objects in the peer to peer system with respect to the objective function.

Identify Relevant Peers. ADiT only distributes the top-k queries to *relevant* peers. A peer p is relevant iff the last delivered object of peer p (i.e. the one with the maximum remaining score on peer p) is among the best k objects of already fetched objects, otherwise peer p is irrelevant and can be pruned, since peer p cannot return a better object than its last published object. The set of relevant peers is updated in each iteration.

Calculating Individual Fetch Sizes. In each iteration ADiT assigns an individual fetch size k'_p to each relevant peer p. The fetch size is determined using the enhanced heuristics discussed in section 3.

```
1  program ADiT (IN string tableName, IN Number k,
2                IN Set<Restriction> restr,
3                IN Set<Function> Sim, IN Function Obj,
4                I_O Map<Number, object> objects)
5     var maxRemScore: Number;
6     var fetchedObjs: Map<Number, object>;
7     var ObjPublished: Number;
8     var relPeers: Set<Peer>;
9     var t: Thread;
10 begin
11     loop
12          maxRemScore = 0;
13          GetRelevantPeers(objects, I_O relPeers);
14          CalcFetchSize(k - objects.count, I_O relPeers);
15         -- broadcast
16          foreach Peer p in relPeers
```

```
17              t = new Thread();
18              t.start(LocalTopKCall(tableName, restr, Sim, Obj,
19                                    p, fetchedObjs, maxRemScore));
20          end-for;
21          -- publish
22          PublishObjects(I_O fetchedObjs, I_O maxRemScore, k,
23                         relPeers, ObjPublished, objects);
24
25      until ObjPublished == k
26  end.
```

Listing 1.1: Pseudo-code for ADiT

Broadcasting Top-K Query. Within each iteration ADiT gathers objects to satisfy the distributed top-k query. Therefore, ADiT distributes the query throughout the system and obtains k'_p objects from each peer *in parallel*.

For each relevant peer ADiT starts a separate thread (*LocalTopKCall*) which encapsulates two major tasks: (1) execution of a local top-k query and (2) updating of the maximum remaining score if it changed (1.2).

```
1  program LocalTopKCall (IN string tableName,
2                         IN Set<Restriction> restr,
3                         IN Set<Function> Sim,
4                         IN Function Obj, IN Peer p,
5                         I_O Map<Number, object> fetchedObjs,
6                         I_O Number maxRemScore)
7  begin
8
9      p.TQQA(tableName, q = 0, p.k', restr,
10             searchType = AT_MOST,
11             Sim, Obj, p.Objects));
12
13     lock(fetchedObjs, maxRemScore);
14         fetchedObjs.AddAll(p.Objects);
15
16         if p.maxScore > maxRemScore then
17            maxRemScore = p.maxScore;
18         end-if;
19     end-lock;
20 end.
```

Listing 1.2: Pseudo-code for sending a top-k query to a certain peer

The first part shows the call of a local *TQQA* query processor [7] which is reentrant, i.e. gathering k'=5 objects in the first iteration and k'=10 objects in the second iteration finally gives the best 15 objects from peer p. After the best (or even next in each following iteration) k' objects have been retrieved they are added to a global buffer. Finally the maximum remaining score is updated in case peer p has a higher maximum score than all other peers.

Publishing Objects. The last step in each iteration is the publishing of relevant objects (Listing 1.3). Since ADiT is an *exact distributed top-k query*

processing approach it is necessary to wait for all peers to return at least one result. This is indicated with the *waitForAll* method. After all peers provided their results ADiT iterates over the *sorted* map and tests for each object whether its score is greater or equal than the maximum remaining score. In that case an object can be published. ADiT stops when enough objects have been published.

```
1  program PublishObjects (I_O Map<Number, object> fetchedObjs,
2                          I_O Number maxRemScore, IN Number k,
3                          IN Set<Peer> relPeers,
4                          I_O Map<Number, object> objects)
5  begin
6      waitForAll(relPeers);
7
8      foreach Element e in fetchedObjs.Elements
9          if e.Score >= maxRemScore
10             objects[e.Score] = e.Object;
11
12             if objects.count == k then
13                 break;
14             end-if
15         else
16             break;
17         end-if
18     end-for
19 end.
```

Listing 1.3: Pseudo-code for the publishing of objects in ADiT

4 Prototype and Experiments

ADiT has been completely implemented in PL-SQL [14,13] as a set of stored procedures [22]. To compare ADiT against a state of the art distributed top-k query processing technique we also implemented the *algorithm with remainder top-k queries (short ARTO) [24]* in this database layer.

4.1 Experimental Setup

We performed experiments on 2 databases: One filled with randomly generated data, and the other consisting of a single relation containing 68 categorical attributes taken from the *UCI Machine Learning Repository* [1,15] which contains over 2.400.000 entries in this single relation which we distributed among the peers in the network such that the size of the database of each peer varied between 5.000 objects and 500.000 objects.

Within this section we present various diagrams generated from the data produced by the conducted test runs. We primarily focused on the *system effort* caused by a certain query and on the *query response time*. To make precise statements about ADiT and the enhanced heuristics we used the basic heuristics with

a $consFactor$ of 2 and four other heuristics to compare them to the enhanced heuristics:

1. $k'_p = k$
2. $k'_p = 1$
3. $k'_p = \lceil \frac{k}{N} \rceil$
4. $k'_p = \lfloor \frac{k}{N} \rfloor$
5. $k'_p = min(k, 2 * \lceil \frac{N_{Size}}{k} \rceil * \frac{k}{N_{Size}}))$

For an easier comparison of the achieved results we defined two ratios: gain with respect to system effort is defined in equation 12; gain achieved for the query response time is shown in equation 13. The respective ratios for the comparison with ARTO are defined accordingly.

$$Ratio_{SE} = \frac{SystemEffort_{heuristic_i}}{SystemEffort_{heuristic_{enhanced}}} \tag{12}$$

$$Ratio_{QAT} = \frac{QueryAnswerTime_{heuristic_i}}{QueryAnswerTime_{heuristic_{enhanced}}} \tag{13}$$

4.2 Discussion of Results

In figure 1 we can see the $Ratio_{QAT}$ for a query with 4 restrictions. Comparing with figure 3 we can see that all curves get higher in a peer to peer network with 49 peers. Additionally, these first figures already show that the heuristics $k'_p = 1$ is not a good choice since it involves high interaction between the query initiator and the other peers. We can also observe that for the query response time the gain over ARTO is rapidly increasing when the amount of searched objects increases. The ratio is growing fast because ARTO needs more sequential message processing when the search amount increases (when the first parallel call was not sufficient).

In figure 2 and figure 3 we can see $Ratio_{SE}$ and $Ratio_{QAT}$ for a query with 4 restrictions in a peer to peer network with 49 peers storing census data. We can observe that the ratio $Ratio_{SE}$ and $Ratio_{QAT}$ are almost identical with respect to their curves. They only differ in the magnitude which is a little higher for the $Ratio_{QAT}$. This means that the usage of ADiT brings slightly more benefits to a single user than to the whole peer to peer system. This result can be observed over all of the tests. The reason for this behaviour is that ADiT tries to fetch fewer objects from less important peers. Thus these peers do not influence the search process that much than in a setting where all peers are contributing the same amount of objects. Another reason is that the search time is dominated by the slowest peer. Avoiding high interaction and fetching few objects from such peers can clearly boost query processing. Another observation is that ARTO has a lower system effort for a small search amount. This can be seen in figure 2. The

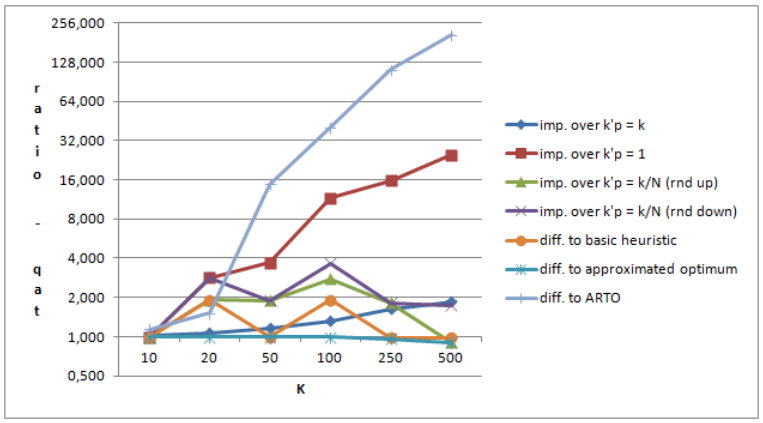

Fig. 1: Ratio for query response time between *enhanced heuristics*, approximated optimum, ARTO and five different approaches to determine the fetch size k'_p in a peer to peer system with 19 peers and varying search amount K and 4 restrictions on census data.

reason for this is that ARTO can answer queries with fewer messages and fewer transmitted objects. This is because ARTO sequentially asks the peer with the highest remaining score for further objects which results in fewer work for the remaining peers. However in figure 3 we can observe that the query response time is better for ADiT in the same scenario.

In figure 4 and figure 5 we can see the $Ratio_{SE}$ and the $Ratio_{QAT}$ for a query with 12 restrictions in a peer to peer network with 49 peers storing census data. In these two figures we can observe the situation where the enhanced heuristics needs more iterations than the heuristics fetching $k'_p = k$ objects. This situation only occurred once in all of the test cases. Additionally, we see the same effect as in figure 2 and figure 3, i.e. the curves are very similar but the $Ratio_{QAT}$ is a little higher than $Ratio_{SE}$. When comparing figures 4 and 5 with figures 2 and 3 we can observe that the magnitude of the ratios is almost independent of the amount of restrictions. Furthermore, we observe in figure 4 and figure 5 that the ratios for ARTO increases at the point where the search amount exceeds the amount of peers in the network. This shows that it is better to ask each peer for more than only one object even when calling them sequentially.

The most important observations gathered through the performed test runs on random and US Census data are:

1. ADiT is up to 200 times faster than ARTO in case the search amount gets higher than the amount of peers in the network.
2. The system effort caused by ADiT is up to 8 times lower than the system effort caused by ARTO in case the search amount gets higher than the number of peers in the network.
3. The query response time of ARTO is in most cases worse than the query response time achieved with any of the presented ADiT heuristics.

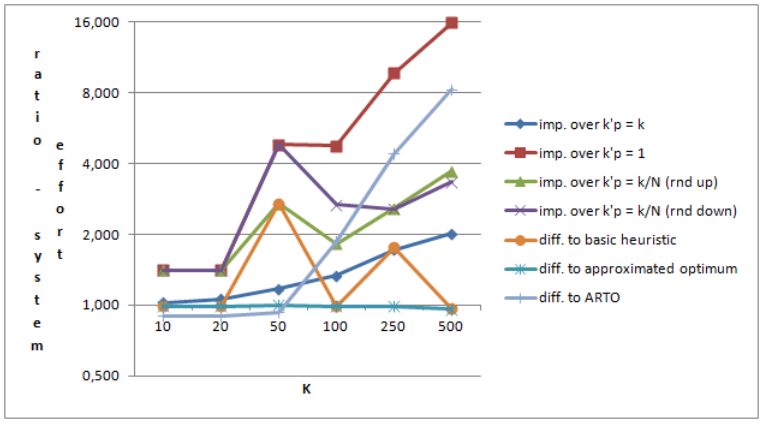

Fig. 2: Ratio for system effort between *enhanced heuristics*, approximated optimum, ARTO and five different approaches to determine the fetch size k'_p in a peer to peer system with 49 peers and varying search amount K and 4 restrictions on census data.

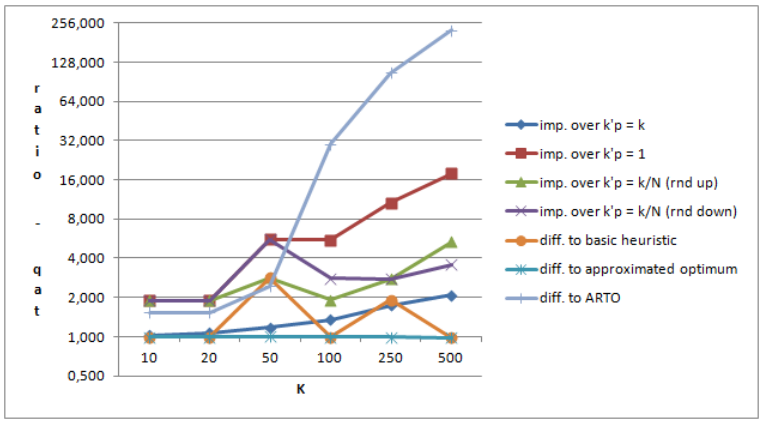

Fig. 3: Ratio for query response time between *enhanced heuristics*, approximated optimum, ARTO and five different approaches to determine the fetch size k'_p in a peer to peer system with 49 peers and varying search amount K and 4 restrictions on census data.

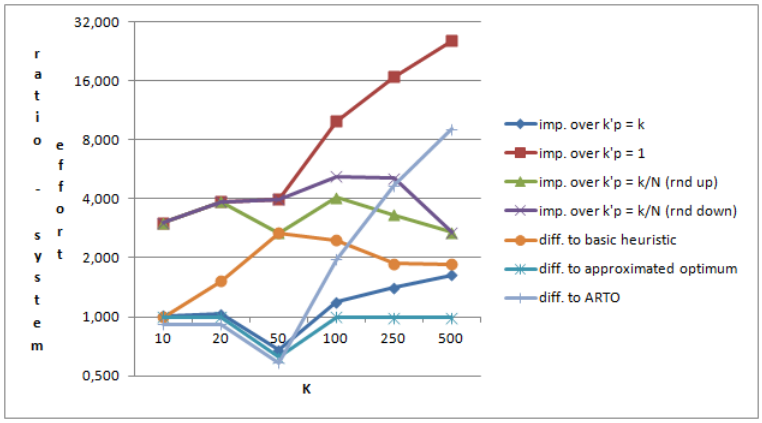

Fig. 4: Ratio for system effort between *enhanced heuristics*, approximated optimum, ARTO and five different approaches to determine the fetch size k'_p in a peer to peer system with 49 peers and varying search amount K and 12 restrictions on census data.

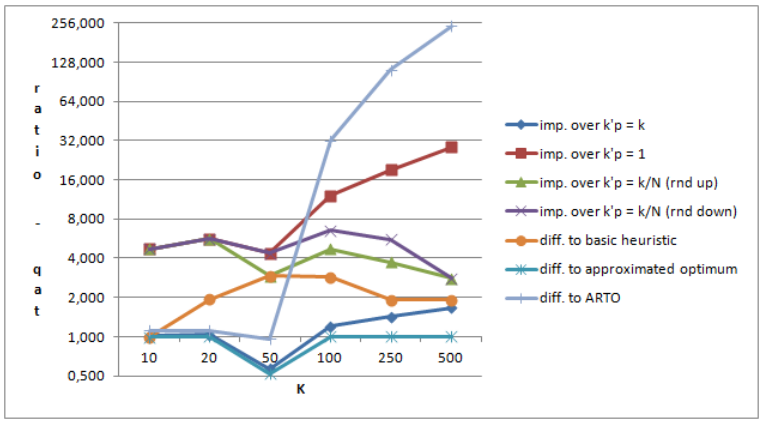

Fig. 5: Ratio for query response time between *enhanced heuristics*, approximated optimum, ARTO and five different approaches to determine the fetch size k'_p in a peer to peer system with 49 peers and varying search amount K and 12 restrictions on census data.

Additionally, we found some characteristics appearing in almost all test runs:

- The enhanced heuristics is close to the approximated optimum gathered through the extensive tests on the *US Census Data (1990) Data Set*.
- The enhanced heuristics is better than all other presented heuristics, except in one single query (see figure 4 and figure 5).
- The enhanced heuristics is between 2 and 32 times faster than the heuristics always fetching *1 object* from each peer in parallel.
- The enhanced heuristics is about 3 to 8 times faster than heuristics fetching $\lceil \frac{k}{N} \rceil$ or $\lfloor \frac{k}{N} \rfloor$ objects from each peer in parallel.
- The enhanced heuristics is between 1.5 and 2.5 times faster than the heuristics fetching k *objects* from each peer in parallel.
- The basic heuristics and the heuristics fetching k *objects* from each peer in parallel turned out to be better than the other heuristics.

5 Conclusion

We discussed distributed top-k query processing from a new perspective. We motivated the need for an adaptive distributed top-k query processing approach (short ADiT) and defined two goal measures, namely (1) the system effort and (2) the query response time. Based on data gathered through extensive experiments we derived a heuristics which can be used to determine a separate fetch size for each peer. We tested the developed heuristics with a large real data set, namely the US Census Data. In these tests we compared the enhanced heuristics against other heuristics and against ARTO [24]. We could show that ADiT can accelerate the query response time and reduce the consumption of system resources significantly. Furthermore, we saw that the *enhanced heuristics* is in most cases close to the best system effort and query response time approximately determined upfront. Additionally, we found that a heuristics fetching more objects is usually the better choice since searching and transmitting a few more objects has much lower costs than sending an additional request. Last but not least the gains achieved with ADiT increase with the size of the peer to peer network and the number of requested results k.

References

1. UCI Machine Learning Repository, US Census Data 1990 `http://archive.ics.uci.edu/ml/datasets/US+Census+Data+(1990)`, 2012.
2. R. Akbarinia, E. Pacitti, and P. Valduriez. Reducing network traffic in unstructured P2P systems using Top-k queries. *Distrib. Parallel Databases*, 19:67–86, 2006.
3. R. Akbarinia, E. Pacitti, and P. Valduriez. Best position algorithms for top-k queries. In *Proc. VLDB '07, 2007.*
4. W.-T. Balke, W. Nejdl, W. Siberski, and U. Thaden. Progressive Distributed Top-k Retrieval in Peer-to-Peer Networks. In *Proc. ICDE '05, IEEE CS, 2005.*

5. N. Bruno, S. Chaudhuri, and L. Gravano. Top-k selection queries over relational databases: Mapping strategies and performance evaluation. *ACM Trans. Database Syst.*, 27:153–187, 2002.
6. W. Conner, S.-w. Hwang, and K. Nahrstedt. Unified Framework for Top-k Query Processing in Peer-to-Peer Networks. 2007.
7. C. Dabringer. *Efficient Local and Distributed Query Processing in a Biomedical Environment*. PhD thesis, Alpen Adria Universität Klagenfurt, 2012.
8. C. Dabringer and J. Eder. Efficient top-k retrieval for user preference queries. In *Proc. ACM Symposium on Applied Computing*, SAC '11, ACM, 2011.
9. C. Dabringer and J. Eder. Fast top-k query answering. In *Proc. 22nd int. conf. on Database and expert systems appl. -Part II*, DEXA'11, Springer-Verlag, 2011.
10. J. Eder, C. Dabringer, M. Schicho, and K. Stark. *Information Systems for Federated Biobanks*, TLDKS Vol 1(1), Springer-Verlag, 2009.
11. Q. Fang and G. Yang. Efficient top-k query processing algorithms in highly distributed environments. *Journal of Computers*, 9(9):2000–2006, 2014.
12. Q. Fang, Y. Zhao, G. Yang, B. Wang, and W. Zheng. Best Position Algorithms for Top-k Query Processing in Highly Distributed Environments. In *Proc 1st Int. Conf. Networking and Distributed Computing*, ICNDC '10, IEEE Computer Society, 2010.
13. S. Feuerstein. *Oracle PL/SQL Programming Guide to Oracle8i Features*. O'Reilly & Associates, 1999.
14. S. Feuerstein and B. Pribyl. *Oracle PL/SQL programming*. 2009.
15. A. Frank and A. Asuncion. UCI Machine Learning Repository, 2010.
16. H. Frank and J. Eder. Towards an automatic integration of statecharts. In *Proc. ER'99, 1999*.
17. R. Hagihara, M. Shinohara, T. Hara, and S. Nishio. A Message Processing Method for Top-k Query for Traffic Reduction in Ad Hoc Networks. In *Proc. Int. Conf. on Mobile Data Management*, MDM '09, IEEE Computer Society, 2009.
18. V. Hristidis and Y. Papakonstantinou. Algorithms and applications for answering ranked queries using ranked views. *The VLDB Journal*, 13(1):49–70, 2004.
19. M. Hua, J. Pei, A. W. C. Fu, X. Lin, and H.-F. Leung. Efficiently answering top-k typicality queries on large databases. In *Proc. VLDB* 2007.
20. I. F. Ilyas, G. Beskales, and M. A. Soliman. A survey of top-k query processing techniques in relational database systems. *ACM Comput. Surv.*, 40(4):1–58, 2008.
21. N. Mamoulis, M. L. Yiu, K. H. Cheng, and D. W. Cheung. Efficient top-k aggregation of ranked inputs. *ACM Trans. Database Syst.*, 32(3):19, 2007.
22. K. T. Owens. *Building intelligent databases with Oracle PL/SQL, Triggers, and stored procedures (2nd ed.)*. Prentice-Hall, 1998.
23. C. Re, N. Dalvi, and D. Suciu. Efficient Top-k Query Evaluation on Probabilistic Data. *Data Engineering, 2007.*
24. N. H. Ryeng, A. Vlachou, C. Doulkeridis, and K. Nørvåg. Efficient distributed top-k query processing with caching. In *Proc. 16th Int. Conf. on Database systems for advanced applications: DASFAA'11, Springer-Verlag, 2011.*
25. A. Vlachou, C. Doulkeridis, K. Nørvåg, and M. Vazirgiannis. On efficient top-k query processing in highly distributed environments. In *Proc. 2008 ACM SIGMOD Int. Conf. on Management of Data*, ACM press, 2008.
26. Z. Zhang, S.-w. Hwang, K. C.-C. Chang, M. Wang, C. A. Lang, and Y.-c. Chang. Boolean + ranking: querying a database by k-constrained optimization. In *Proc. 2006 ACM SIGMOD Int. Conf. on Management of Data*, ACM press, 2006.